
ROADRUNNER
ON THE ROAD TO
TRINITY

For almost 50 years, nuclear weapons were tested explosively, first at the Pacific Proving Grounds and until 1992, at the Nevada Test Site. Now nuclear weapons are simulated on supercomputers using codes in which computational physicists have captured the essence of weapon performance.

The simulations are part of Stockpile Stewardship, the National Nuclear Security Administration (NNSA) program established as a means of assessing the safety, security, and reliability of the United States' stockpiled nuclear weapons in the absence of nuclear testing. The assesments drawn from simulations enable the directors of the three U.S. national security laboratories—Los Alamos, Lawrence Livermore, and Sandia—to annually inform the president of the United States, through the secretaries of Energy and Defense and the Nuclear Weapons Council, that the stockpile of current weapons remains, in their view, safe, secure, and effective.

Roadrunner was a step on the road to Trinity, and it was a huge step—one for the record books.

Providing the input that underwrites such an important assesment is a daunting task, so the supercomputers that weapons scientists use are always on computation's leading edge. Los Alamos National Laboratory's Roadrunner supercomputer won world fame for breaking a computing speed record, but Los Alamos is pushing its computing capabilities to newer, higher levels to meet the growing needs of Stockpile Stewardship. So the Laboratory is looking forward to its next big computer, a machine already named but still to be fully conceived: Trinity. Trinity will be designed and procured by the New Mexico Alliance for Computing at Extreme Scales, a partnership between Los Alamos and Sandia. Roadrunner was a step on the road to Trinity, and it was a huge step—one for the record books.

The Need for Speed

Like its namesake, the state bird of New Mexico, Roadrunner was certainly speedy. In June 2008 it took over first place on the Top500, the international list of the world's fastest supercomputers. It was the world's first petascale supercomputer, that is, the first to reach a sustained quadrillion (thousand trillion)—1,000,000,000,000,000—calculations (called "floating-point operations") per second: petaflops.

Roadrunner was first in another way as well. It had a unique design (architecture) for a supercomputer. Cheryl Wampler, the Laboratory's deputy director for the Advanced Simulation and Computing (ASC) program—the NNSA program that oversees the development of supercomputing technology for Stockpile Stewardship—describes Roadrunner as "a historic machine, a pioneer." Then she adds, "It was also controversial."

The controversy was about that unique architecture. Roadrunner was a "hybrid" in that it combined two different kinds of processors (chips that read and carry out program instructions). It had 6,563 dual-core general-purpose processors (AMD Opterons)—the kind found in almost all computers, except that each was actually two processors in one because of its two compute cores. Each core was then linked to a special graphics processor (PowerXCell 8i) called the Cell. The Cell was an enhanced version of a specialized processor originally designed for the Sony PlayStation 3. For Roadrunner the Cell was adapted specifically to support scientific computing.

Although hybrid computers had existed before Roadrunner, no one had ever tried that approach on such a large scale, and many doubted that a hybrid supercomputer would work. So for Los Alamos and IBM, who collaborated to design the computer, Roadrunner was a leap of faith—but a leap with purpose.

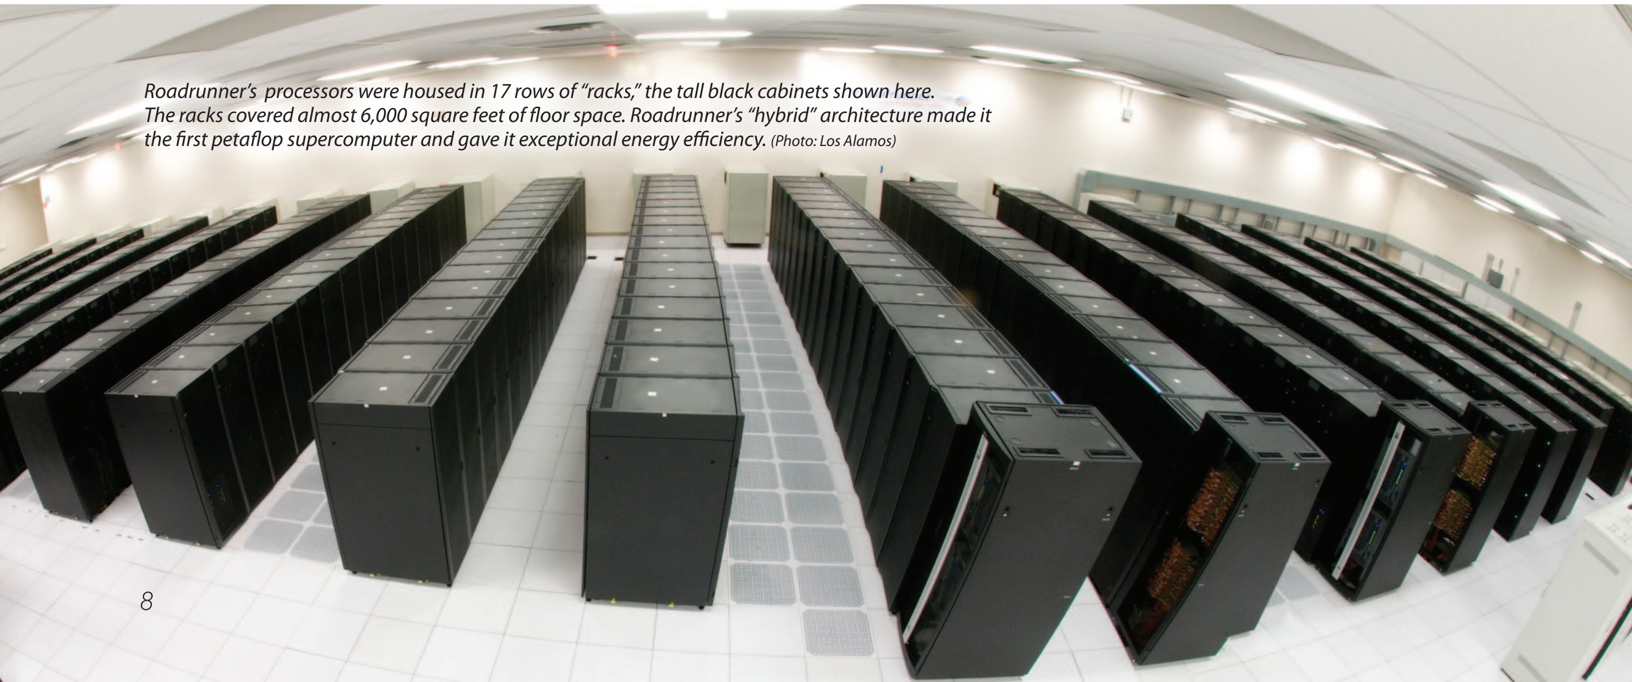

Roadrunner's processors were housed in 17 rows of "racks," the tall black cabinets shown here. The racks covered almost 6,000 square feet of floor space. Roadrunner's "hybrid" architecture made it the first petaflop supercomputer and gave it exceptional energy efficiency. (Photo: Los Alamos)

High-speed calculation was certainly part of the purpose, as Jim Lujan, in the Laboratory's High Performance Computing Division, explains: "We wanted to take the next major step in computing. Computers get faster every year, but we needed an exponential increase to solve the physics problems we have in the NNSA."

The physics problems associated with nuclear weapon simulations are both big and complex, especially since the simulations need to be of an entire weapon, not just its individual components. Only a very fast supercomputer can complete all the calculations involved in a reasonable amount of time.

Computers get faster every year, but we needed an exponential increase to solve the physics problems we have in the NNSA.

"We're looking at all the physics involved with a weapon, from initiation to completion of a detonation," says Lujan. "With slower, less powerful computers, we were only able to do pieces—primaries, secondaries, or parts of the ignition. Now we're looking at the whole system, end-to-end. We could have done it 10 years ago with the computing systems we had then, but it would have taken years upon years upon years, and that doesn't make it a tractable problem."

It's All in the Details

Although simulation only *predicts* weapon performance, when nuclear testing *demonstrated* it, simulation pays a dividend in detail—detail about all the internal processes that feed into performance.

The scientists who develop weapons codes lay the groundwork for that dividend by including equations and data representing all aspects of a weapon: the size and shape of components, the chemical makeup and behavioral tendencies of constituent materials (from plastics to plutonium), and the phenomena acting on everything inside a weapon during detonation. Data collected during the days of nuclear testing inform the code developers, as do new data drawn from experiments on the many materials used in nuclear weapons and on data from *non*nuclear tests of weapon components.

When a computer is fast enough to work all of a code's equations in a reasonable time span, the resulting simulations deepen scientists' understanding of weapon behavior. The faster the computer, the more equations it can handle and, therefore, the more detail its simulations can provide about smaller pieces of the whole.

That small pieces are critical to the fate of an entire system is just as true for nuclear weapons as it was for the space shuttle Challenger on January 18, 1986, when the failure of an O-ring just 0.280 inch in diameter led to the space craft's destruction and the death of everyone aboard.

So scientists use simulation to learn how all parts of a weapon behave during detonation and use what they learn to predict the continued viability of the weapons in the stockpile. The information harvested from each simulation, along with the data from continued experiments and nonnuclear tests, is then used to improve the codes so that future simulations, provided they are run on a fast-enough computer, will be even more detailed and more accurately predictive of performance.

Matching Code to Computer

Roadrunner's speed was derived from its architecture. The two processors shared functions, with the Cell taking on the most computationally intense parts of a calculation to hurry the work along. Essentially, the Cell acted as a computational accelerator. But the division of tasks was not automatic; it was achieved through programming, so preexisting codes had to be revised by their developers. In fact, the code work began years before Roadrunner was developed.

Rewriting a code for Roadrunner required developers to puzzle out the best way to divide calculations between the processors and to carefully consider how to distribute data so that it would be in the right place when needed. That kind of rethinking was a significant challenge, but there were rewards for successful code revision. In work for Stockpile Stewardship, Roadrunner took on a troubling, long-standing gap in understanding about how energy flows in a detonation and how yield is dependent upon that flow. The big computer made a significant contribution to that understanding.

As code developers grappled with the new machine and its unfamiliar architecture, they were hearing a warning shot. They were being alerted to upcoming changes in high-performance computing.

General-science researchers also had time on Roadrunner, during the computer's first six months of work—its "shakedown" period before it was transitioned to classified status. Scientists using the "ML" code used Roadrunner to study the genetic sequences of HIV and to build an HIV family tree that they hope will help in the search for a vaccine for the virus. Other codes were used to produce breakthroughs in fields such as materials science, astronomy, and laser-plasma interaction.

Brian Albright of the Computational Physics Division was one of several researchers using the VPIC (vector particle-in-cell) code for projects that included the laser-plasma interaction studies. He noted that adjusting the code for Roadrunner had far-reaching benefits. “We set up several code teams to adapt VPIC for Roadrunner, and the changes made VPIC a more general code, not tied to a single computer architecture.”

A Shot across the Bow

As demonstrated by successes with VPIC, code work done to accommodate Roadrunner will be generically applicable to emerging new types of supercomputers. That fact highlights another significant purpose behind Roadrunner: demonstrating that the weapons codes need to be broken free of their dependence on computer architectures that are rapidly becoming antiquated. As code developers grappled with the new machine and its unfamiliar architecture, they were hearing a warning shot. They were being alerted to upcoming changes in high-performance computing—changes that everyone working with Stockpile Stewardship is now embracing.

As Wampler explains it, “Roadrunner was an ‘Advanced Architecture’ supercomputer, a category defined by the ASC program. Advanced Architecture machines are meant to evolve computing into the future.”

As expected with any ahead-of-the-curve technology, Roadrunner could seem to be, as Wampler describes it, “a little flaky,” but that was because of its evolutionary nature; it was not necessarily intended to be used, or used comfortably, by all codes. However, Roadrunner did exactly what it *was* intended to do regarding the weapons codes: it got the codes moving toward new architectures. The codes will *have* to run on new architectures . . . and soon. Roadrunner was challenging because the supercomputing future will be challenging.

Cielo—End of an Era

Things are already changing in the Laboratory’s Nicholas C. Metropolis Center for Modeling and Simulation (also known as the Strategic Computing Center, or SCC), where Roadrunner was installed. Roadrunner’s time has run out, and it is going away. The big machine will soon be dismantled and will assume its rightful place in history as the focal point of dramatic change at the start of a new era in computing history.

What remains behind at the SCC is an equally powerful supercomputer named Cielo, which will be active for the next few years and is already hard at work; it has been overlapping with Roadrunner since January 2011. Cielo is more like previous Stockpile Stewardship workhorses (not a hybrid) while operating at Roadrunner scale: 1.43 petaflops. It was built to accommodate the present-day weapons codes—no

An explosion containment vessel is shown being prepared for a test at the Laboratory’s Dual-Axis Radiographic Hydrotest (DARHT) facility. DARHT is one of several Los Alamos facilities where tests and experiments generate data that Los Alamos scientists incorporate into codes that are run on supercomputers like Roadrunner for simulations of full-scale nuclear weapon tests. (Photo: Los Alamos)

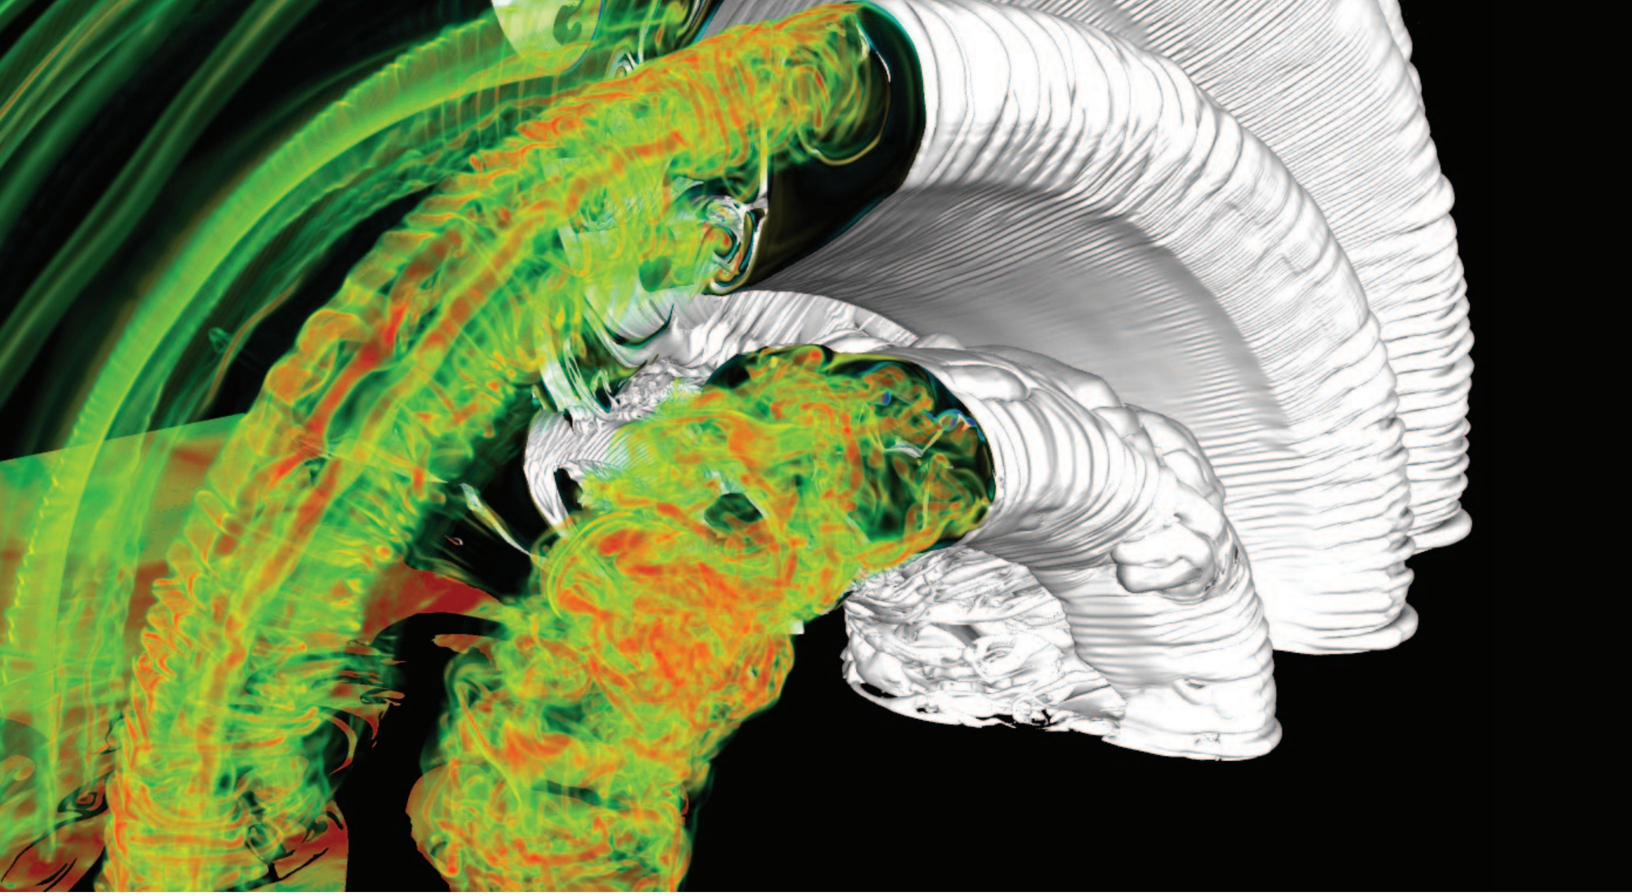

At Los Alamos, simulations such as those run on Roadrunner and currently on Cielo can be seen as full-color 3D visualizations such as this one showing the growth of turbulence in the gas inside an inertial confinement fusion capsule. This image is a time shot from a full simulation done on Cielo. (Photo: Los Alamos)

extensive adaptations needed—and its continuity with previous systems will give *all* the codes time to be prepared for future machines.

After Cielo, we won't be able to say, "Give me another one just like the last one." High-performance computing is moving in a new direction, and we have to move with it.

The future is coming quickly because Cielo will be on its way out by 2015. And when it is gone, it will *really* be gone. Cielo is the last of its kind.

"After Cielo," says Wampler, "we won't be able to say, 'Give me another one just like the last one.' High-performance computing is moving in a new direction, and we have to move with it."

Keeping up is essential because the national laboratories do not typically have their supercomputer's hardware built for them from scratch. Instead, their computers are built with commercially available technology as the starting point. That approach is more economical. Besides, the most rapid changes in computer technology are occurring in the private sector; the national laboratories must track commercial technology trends if their supercomputers are to be on the cutting edge, easily maintained, and able to minimize downtime.

Market trends are driving changes, with companies trying new architectures and developing new technologies such as different kinds of processors, memory units, and interconnects, and all those innovations add up to an entirely new software environment.

But the needs of Stockpile Stewardship are also driving change. The program's demands are growing and require greater and greater computational power because nuclear testing, and the certainty it provided about weapon reliability, is receding further and further into the past. Indeed the stockpiled weapons are not the same as they were when nuclear tests produced as much information as we were able to capture. The weapons' components, materials, and systems have aged, possibly changing in ways that make them behave differently. Some of the components have been remanufactured. As complex as modern nuclear weapons already are, they become even more complex as they age, so time is raising new questions that are waiting for answers.

Trinity is expected to be the first platform large enough and fast enough to begin to accommodate finely resolved 3D calculations for full-scale, end-to-end weapons calculations.

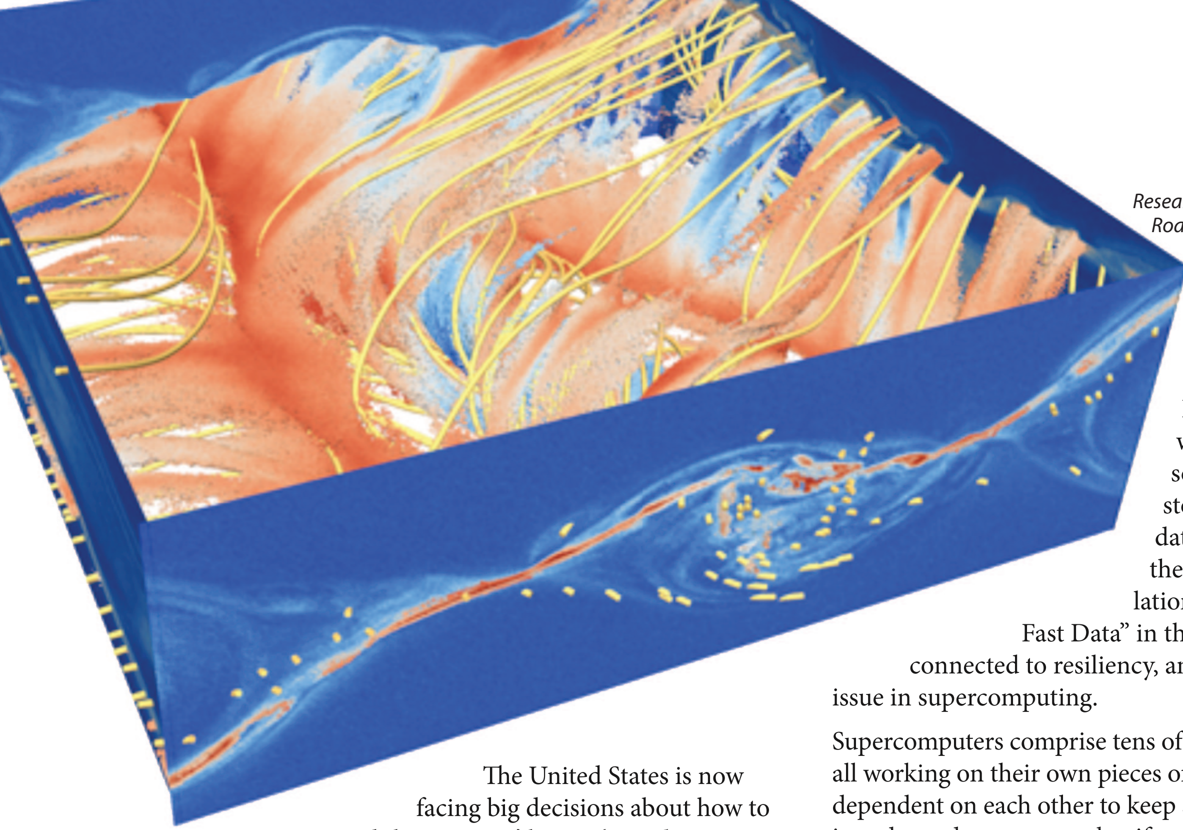

Researchers used the VPIC code on Roadrunner to produce this simulation of magnetic reconnection. VPIC was also used to simulate laser-plasma interactions.

The United States is now facing big decisions about how to extend the weapons' lives. Those decisions will depend on how much scientists know about what makes a weapon work . . . or not work. As a result, the national security labs need even bigger and more detailed simulations with superior resolution, in 3D, for capturing fine-scale features. And the simulations will need to cover all the physics associated with a full-scale weapon.

In other words, the laboratories will need simulations that are vastly more predictive than what today's supercomputers can provide. So the national security science labs have to take dead aim on the supercomputing future.

Forward to Trinity

Eventually, maybe as soon as 2020, supercomputers will reach exascale—one quintillion (1,000,000,000,000,000,000) calculations per second—making them 1,000 times faster than Roadrunner. Such speed bodes well for the needs of the national security labs. But along with being fast, future supercomputers will need to be energy efficient if they are to be welcome in an increasingly energy-conscious world.

Bigger systems have many, many more processors, each of which can fail, causing the whole machine to stop.

Roadrunner made great strides in energy efficiency. Although bigger and older than Cielo, Roadrunner used significantly less energy to achieve essentially the same speed. Future supercomputers will need to emulate that success . . . and improve on it if speed and power are to continue increasing.

Future supercomputers will also need new memory solutions for handling and storing the vast amounts of data used in simulations and the "big data" those same simulations generate (see "Big Data, Fast Data" in this issue). Memory is closely connected to resiliency, and resiliency is a very big issue in supercomputing.

Supercomputers comprise tens of thousands of processors all working on their own pieces of a calculation and all dependent on each other to keep a calculation going. That interdependence means that if one processor fails, they all have to stop. And with more and more processors coming into play as systems grow, the number of failures is growing as well.

Supercomputing codes mitigate the losses from such failures by "checkpointing," halting at regular intervals to save the current state of a calculation to storage (long-term memory). Then when a failure occurs, the latest checkpoint can be recalled and the calculation restarted at that point. The faster a checkpoint can be saved and the faster it can be recalled after a failure, the less downtime there will be—and that is resiliency. The designers of tomorrow's supercomputers will need to look hard at ways to make them as resilient as possible.

Getting to better efficiency and resiliency, and reaching exascale, will not happen all at once. It is happening in stages, and the next stage of the journey is Los Alamos' next addition: Trinity.

Roadrunner got everyone thinking in new ways about how to build a supercomputer.

Trinity will be an Advanced Technology System—a new ASC category replacing Advanced Architecture. This new machine, tentatively projected for installation in 2015–2016, could be 40 to 50 times faster than Roadrunner and Cielo. Although it is expected to be, like Roadrunner, a significant break with the past, it will have to serve, like Cielo, as the working platform for *all* the Stockpile Stewardship codes that need to run at large scales. As such, Trinity is expected to be the first platform large enough and fast enough to begin to accommodate finely resolved 3D calculations for full-scale, end-to-end weapons calculations.

Exactly what Trinity will be like is still under discussion. Trinity's designers are thinking in particular about what Gary Grider, of the Laboratory's High Performance Computing Division, calls "being smarter about the use of silicon"—silicon being the key material that makes up the transistors on the computer's processors (chips).

When Roadrunner's specialized processor, the Cell, assumed some functions that the Opteron processors would otherwise have done, it was taking over operations that were not only computationally complex but also computationally *expensive*—requiring a great deal of time and power and many floating-point operations per second.

Using the Cell for complex and expensive functions was Roadrunner's example of using silicon "smartly." If all a computer's functions are performed on a general-purpose processor (like the Opteron), there will be times when only parts of the processor are in use, even though all parts of the processor, even the inactive ones, are being energized. That adds up to wasted energy, a "dumb" use of silicon.

So Roadrunner's division of labor enabled not only greater speed but also greater energy efficiency. It was a strategy that landed Roadrunner on the Green500 list, a measure of world supercomputer efficiency, in addition to being on the speedy Top 500 list.

"Roadrunner wasn't the first time anyone had thought of using specialized processors that way," says Grider, "but it was the first time anyone had done a really big demonstration of what that means. So we produced a petaflop machine using 3 megawatts of power. At the same time, the Jaguar supercomputer at Oak Ridge used 8 megawatts of power to get a petaflop because they were doing things the old-fashioned way—using only general-purpose processors."

Another old-fashioned way to use silicon involves the movement of data to and from processors. Says Grider, "Computer programs today are written on the model that you move data to the processors, you do calculations, and then you write the results out to memory [storage]. Then later you move data back out of memory to be returned to the processors for new calculations. We can't afford to do that anymore because moving data around requires a lot of power—it's not very efficient. In addition, the time it takes to move data is wasted. It's time spent *not* computing."

There is more than one possible solution. Why not have some processing done right there in memory? Or, why not store data closer to processors? The Laboratory's computer scientists began to think of both possibilities as Roadrunner's huge computations raised questions about efficient versus inefficient data movement.

The Smart Supercomputing Future

Roadrunner got everyone thinking in new ways about how to build a supercomputer. Specialized processors are already being included; Lawrence Livermore's Sequoia, the fastest supercomputer in the world in June 2012, and Oak Ridge's Titan, the current fastest, are examples of that. "So our demonstration with Roadrunner," Grider concludes, "caused everyone to pay attention."

Trinity is on the way to exascale —and who knows what else lies beyond.

Certainly Trinity's designers are paying attention. They have put a lot of new ideas on the table, not only a greater use of specialized processors but also the creation of layered memory, with a "flash" memory to improve efficiency in data movement between processors. There are a lot of decisions to be made before the machine becomes a reality, and the final picture is still out of focus.

But if the exact nature of Trinity is still uncertain, what *is* certain is that it will not do what Cielo does—provide a comfortable environment for current, unadapted weapons codes. So the code developers at Los Alamos and at all the national security labs are already working to radically change their codes to fit the radical architectures to come.

Trinity is on the way. Exascale—and who knows what else—lies beyond. Roadrunner's warning came just in time for everyone to get ready.

~Eileen Patterson